Great Expectations.
Part I: On the Customizability of Generalized Expected Utility*

Francis C. Chu   Joseph Y. Halpern

Department of Computer Science
Cornell University
Ithaca, NY 14853, U.S.A.
Email: {fcc,halpern}@cs.cornell.edu

Abstract

We propose a generalization of expected utility that we call *generalized EU* (GEU), where a decision maker's beliefs are represented by plausibility measures and the decision maker's tastes are represented by general (i.e., not necessarily real-valued) utility functions. We show that every agent, "rational" or not, can be modeled as a GEU maximizer. We then show that we can customize GEU by selectively imposing just the constraints we want. In particular, we show how each of Savage's postulates corresponds to constraints on GEU.

1 Introduction

Many decision rules have been proposed in the literature. Perhaps the best-known approach is based on maximizing expected utility (EU), calculated either with respect to a given (objective) probability measure, as done originally by von Neumann and Morgenstern [1947], or with respect to a probability measure constructed from a preference order on alternatives that satisfies certain postulates, as done originally by Savage [1954]. All these approaches follow the same pattern: they formalize the set of alternatives among which the decision maker (DM) must choose (typically as *acts* or *lotteries*[1]). They then give a set of assumptions (often called *postulates* or *axioms*) such that the DM's preferences on the alternatives satisfy these assumptions iff the preferences have an EU representation, where an *EU representation* of a preference relation is basically a utility function (and a probability measure when acts are involved) such that the relation among the alternatives based on expected utility agrees with the preference relation. Moreover, they show that the representation is essentially unique in that, given two representations of the same preference relation, the utility functions are positive affine transformations of one another and the probability measures are equal. Thus, if the preferences of a DM satisfy the assumptions, then she is behaving *as if* she has quantified her tastes via a real-valued utility function (and her beliefs via a probability measure) and she is relating the alternatives according to their expected utility. The assumptions are typically regarded as criteria for rational behavior, so these results also suggest

*Work supported in part by NSF under grants IIS-0090145 and CTC-0208535 and by the DoD Multidisciplinary University Research Initiative (MURI) program administered by ONR under grant N00014-01-1-0795.

[1]Formally, given a set S of states of the world and another set C of *consequences*, an *act* a is a function from S to C that, intuitively, associates with each state s the consequence of performing a in s. A *lottery* is a probability distribution over consequences; intuitively, the distribution quantifies how likely it is that each consequence occurs.

that if a DM's beliefs are actually described by a probability measure and her tastes are described by a utility function, then she should relate the alternatives according to their expected utility (if she wishes to appear rational).

Despite the appeal of EU maximization, it is well known that people do not follow its tenets in general [Resnik 1987]. As a result, a host of extensions of EU have been proposed that accommodate some of the more systematic violations (see, for example, [Gul 1991; Gilboa and Schmeidler 1989; Giang and Shenoy 2001; Kahneman and Tversky 1979; Luce 2000; Quiggin 1993; Schmeidler 1989; Tversky and Kahneman 1992; Yaari 1987]). Again, the typical approach in the decision theory literature has been to prove representation theorems. These representation theorems essentially view a decision rule $\mathcal{R}$ as a function that maps tastes (and perhaps beliefs, depending on the rule) to a preference relation on alternatives. The theorem then says that, given a rule $\mathcal{R}$, there is a set $\mathcal{A}_{\mathcal{R}}$ of assumptions about preference orders such that a preference relation $\preceq$ satisfies $\mathcal{A}_{\mathcal{R}}$ iff there exists some tastes and beliefs for the agent such that, given these as inputs, $\mathcal{R}$ returns $\preceq$.

Given this plethora of rules, it would be useful to have a general framework in which to study decision making. The framework should also let us understand the relationship between various decision rules. We provide such a framework in this paper.

The basic idea of our approach is to generalize the notion of expected utility so that it applies in as general a context as possible. To this end, we introduce *expectation domains*, which are structures consisting of

- three (component) domains: a *plausibility domain* P, a *utility domain* U, and a *valuation domain* V,
- two binary operators $\oplus : V \times V \to V$ and $\otimes : P \times U \to V$, which are the analogues of $+$ and $\times$ over the reals, and
- a reflexive binary relation $\precsim$ on V (which generalizes $\leq$).

Intuitively, $\otimes$ combines plausibility values and utility values much the same way that $\times$ combines probability and (real) utility, while $\oplus$ combines the products to form the (generalized) expected utility, in a way analogous to adding products of probabilities and utilities in calculating standard expected utility.

We have three domains because we do not want to require that DMs be able to add or multiply plausibility values or utility values, since these could be qualitative (e.g., plausibility values could be "unlikely", "likely", "very likely", etc., and utility values could be "bad", "good", "better", etc.). In general, we do not assume that $\precsim$ is an order (or even a preorder), since we would like to be able to represent as many preference relations and decision rules as possible.

Once we have an expectation domain, DMs can express their tastes and beliefs using components of the expectation domain. More specifically, the DMs express their beliefs using a *plausibility measure* [Friedman and Halpern 1995], whose range is the plausibility domain of the expectation domain (plausibility measures generalize probability measures and a host of other representations of uncertainty, such as sets of probability measures, Choquet capacities, possibility measures, ranking functions, etc.) and they express their tastes using a utility function whose range is the utility domain of the expectation domain. In an expectation domain, it is possible to define a generalization of expected utility, which we call *generalized EU* (GEU). The GEU of an act is basically the sum (in the sense of $\oplus$) of products (in the sense of $\otimes$) of plausibility values and utility values that generalizes the standard definition of (probabilistic) expected utility over the reals in the obvious way.

We start by proving an analogue of Savage's result with respect to the decision rule (Maximizing) GEU.[2] We show that every preference relations on acts has a GEU representation (even those that do not satisfy any of Savage's postulates), where a *GEU representation* of a preference relation basically consists of an expectation domain E, plausibility measure Pl, and utility function $\mathbf{u}$, such that the way acts are related according to their GEU agrees with the preference relation (Theorem 3.1). In other words, no matter what the DM's preference relation on acts is, she behaves as if she has quantified her beliefs by a plausibility measure and her tastes via a utility function, and is relating the acts according to their (generalized) expected utility as defined by the $\oplus$ and $\otimes$ of some expectation domain. That is, we can model *any* DM using GEU, whether or not the DM satisfies any rationality assumptions. An important difference between our result and that of Savage is that he was constructing *EU* representations, which consists of a *real-valued* utility function $\mathbf{u}$ and a *probability* measure Pr (and the expectation domain is fixed, so $\oplus$, $\otimes$, and $\precsim$ are just $+$, $\times$, and $\leq$, respectively).

Given that GEU can represent *all* preference relations, it might be argued that GEU is too general—it offers no guidelines as to how to make decisions. We view this as a feature, not a bug, since our goal is to provide a general framework in which to express and study decision rules, instead of proposing yet another decision rule. Thus the absence of "guidelines" is in fact an absence of *limitations*: we do not want to exclude any possibilities at the outset, even preference relations that are not transitive or are incomplete. From the point of view of a behavioral scientist, this has the advantage of allowing us to represent the preference relations that actually arise in real life, which typically do not satisfy many of the standard assumptions made by decision theorists, and doing so in a potentially compact way (by specifying $\oplus$, $\otimes$, $\precsim$, a plausibility measure Pl, and a utility function $\mathbf{u}$). Perhaps more interesting is that, starting from a framework in which we can represent all preference relations, we can then consider what preference relations have "special" representations, in the sense that the expectation domain, plausibility measure, and utility function in the representation satisfy some (joint) properties. This allows us to show how properties of expectation domains correspond to properties of preference relations. We can then "customize" GEU by placing just the constraints we want. We illustrate this by showing how each of Savage's postulates corresponds in a precise sense to an axiom on GEU.

This ability to customize GEU may be of more interest to computer scientists than to behavioral scientists. If we try to design software agents that make decisions on our behalf, it may not be appropriate to assume that they will represent beliefs using probability measures and tastes using real-valued utilities. For example, the information that a system can obtain may be better modeled by a set of probability measures than a single probability measure, and a user may represent his or her tastes more qualitatively, using words like "terrific" and "terrible", rather than numerically. Using GEU, it should be possible to design agents that make decisions based on more general representations of beliefs and tastes, and customize them so that the the decision-making process satisfies certain "rationality" postulates.

There is yet another advantage of this approach, which is the focus of [Chu and Halpern 2003]. Intuitively, a decision rule maps tastes (and beliefs) to preference relations on acts. Given two decision rules $\mathcal{R}_1$ and $\mathcal{R}_2$, an $\mathcal{R}_1$ *representation of* $\mathcal{R}_2$ is basically a function τ that maps inputs of $\mathcal{R}_2$ to inputs of $\mathcal{R}_1$ that represent the same tastes and beliefs, with the property that $\mathcal{R}_1(\tau(x)) = \mathcal{R}_2(x)$. Thus, τ models, in a precise sense, a user of $\mathcal{R}_2$ as a user of $\mathcal{R}_1$, since τ preserves tastes (and beliefs).

[2] Many decision rules involve optimizing (i.e., maximizing or minimizing) some value function on the acts. Sometimes it is explicitly mentioned whether the function is to be maximized or minimized (e.g., "Minimax Regret" says explicitly to "minimize the maximum regret") while other times only the function name is mentioned and it is implicitly understood what is meant (e.g., "EU" means "maximize EU"). In this paper we use "Maximizing GEU" and "GEU" interchangeably.

In [Chu and Halpern 2003] we show that many decision rules have GEU representations. Moreover, we show that (almost) every decision rule has an *ordinal* GEU representation, where, rather than x and $\tau(x)$ representing exactly the same tastes (and beliefs), they preserve the relation between tastes (and beliefs), without necessarily preserving the magnitude. For example, if outcome o_1 is preferred to outcome o_2 in x, o_1 is also preferred to o_2 in $\tau(x)$, although the magnitude of preference may be different.

Our claim that only "many" decision rules have a GEU representation may seem inconsistent with our earlier claim that every preference relation has a GEU representation. Representing a preference relation is not the same as representing a decision rule. If we again view a decision rule as a function from tastes (and possibly beliefs) to preference relation on alternatives, then decision rule $\mathcal{R}$ represents a preference relation $\preceq$ if there are some tastes and beliefs such that, with these as input, $\mathcal{R}$ returns $\preceq$. On the other hand, $\mathcal{R}_1$ represents $\mathcal{R}_2$ if, roughly speaking, for *all* possible inputs of tastes (and beliefs), $\mathcal{R}_1$ and $\mathcal{R}_2$ return the same preference relation. That is, $\mathcal{R}_1$ and $\mathcal{R}_2$ act essentially the same way as functions.

Although there has been a great deal of work on decision rules, there has been relatively little work on finding general frameworks for representing decision rules. In particular, there has been no attempt to find a decision rule that can represent all preference relations. There has been work in the fuzzy logic community on finding general notions of integration (which essentially amounts to finding notions of expectation) using generalized notions of $\oplus$ and $\otimes$; see, for example, [Benvenuti and Mesiar 2000]. However, the expectation domain used in this work is (a subset of) the reals; arbitrary expectation domains are not considered. Luce [1990, 2000] also considers general addition-like operations applied to utilities, but his goal is to model joint receipts. Receipts are typically modeled in an arguably inefficient way, as commodity bundles; it seems more natural to deal with them as Luce does, in terms of an abstract binary operation, akin to the operations we have here.

The rest of this paper is organized as follows. We cover some basic definitions in Section 2: plausibility domains, utility domains, expectation domains, decision problems, and GEU. We show that every preference relation on acts has a GEU representation in Section 3. In Section 4, we show that each of Savage's postulates corresponds to an axiom on GEU. We conclude in Section 5. Most proofs are deferred to the appendix.

2 Preliminaries

2.1 Plausibility, Utility, and Expectation Domains

Since one of the goals of this paper is to provide a general framework for all of decision theory, we want to represent the tastes and beliefs of the DMs in as general a framework as possible. In particular, we do not want to force the DMs to linearly preorder all consequences and all events (i.e., subsets of the set of states). To this end, we use plausibility measures to represent the beliefs of the DMs and (generalized) utility functions to represent their tastes.

A *plausibility domain* is a set P, partially ordered by $\preceq_P$ (so $\preceq_P$ is a reflexive, antisymmetric, and transitive relation), with two special elements $\bot_P$ and $\top_P$, such that (We often omit the subscript P in $\bot_P$ and $\top_P$ when it is clear from context.) $\bot_P \preceq_P x \preceq_P \top_P$ for all $x \in P$. Given a set S, a function $\mathrm{Pl} : 2^S \to P$ is a *plausibility measure* iff

Pl1. $\mathrm{Pl}(\emptyset) = \bot$,

Pl2. $\mathrm{Pl}(S) = \top$, and

Pl3. if $X \subseteq Y$ then $\mathrm{Pl}(X) \preceq \mathrm{Pl}(Y)$.

Clearly plausibility measures are generalizations of probability measures. As pointed out in [Friedman and Halpern 1995], plausibility measures generalize a host of other representations of uncertainty as well. Note that while the probability of any two sets must be comparable (since $\mathbb{R}$ is totally ordered), the plausibility of two sets may be incomparable.

We also want to represent the tastes of DMs using something more general than $\mathbb{R}$, so we allow the range of utility functions to be utility domains, where a *utility domain* is a set U endowed with a reflexive binary relation $\precsim_U$. Intuitively, elements of U represent the strength of likes and dislikes of the DM while elements of P represent the strength of her beliefs. Note that we do not require the DM's preference to be transitive (although we can certainly add this requirement). Experimental evidence shows that DM's preferences occasionally do seem to violate transitivity.

Once we have plausibility and utility, we want to combine them to form expected utility. To do this, we introduce expectation domains, which have utility domains, plausibility domains, and operators $\oplus$ (the analogue of $+$) and $\otimes$ (the analogue of $\times$).[3] More formally, an *expectation domain* is a tuple $E = (U, P, V, \oplus, \otimes)$, where $(U, \precsim_U)$ is a utility domain, $(P, \preceq_P)$ is a plausibility domain, $(V, \precsim_V)$ is a valuation domain (where $\precsim_V$ is a reflexive binary relation), $\otimes : P \times U \to V$, and $\oplus : V \times V \to V$. There are four requirements on expectation domains:

- $E1$. $(x \oplus y) \oplus z = x \oplus (y \oplus z)$;
- $E2$. $x \oplus y = y \oplus x$;
- $E3$. $\top \otimes x = x$;
- $E4$. $(U, \precsim_U)$ is a substructure of $(V, \precsim_V)$.

$E1$ and $E2$ say that $\oplus$ is associative and commutative. $E3$ says that $\top$ is the left-identity of $\otimes$ and $E4$ ensures that the expectation domain respects the relation on utility values.

Note that we do not require that $\oplus$ be monotonic; that is, we do not require that for all $x, y, z \in V$,

$$\text{if } x \precsim_V y \text{ then } x \oplus z \precsim_V y \oplus z. \tag{2.1}$$

We say that E is *monotonic* iff (2.1) holds. It turns out that monotonicity does not really make a difference by itself; see Corollary 3.2.

Recall that in the standard case, $\bot = 0$ and $0 \times x$ is the identity for $+$. In general, we do not assume that $\bot \otimes u$ is the identity $\oplus$ (or that $\oplus$ even has an identity). We say that E *has a* $\oplus$ *identity* iff

$$(\bot \otimes u) \oplus x = x \text{ for all } u \in U \text{ and } x \in V. \tag{2.2}$$

Because $\oplus$ is commutative, there can clearly be at most one identity for $\oplus$, so if (2.2) holds, then $\bot \otimes u_1 = \bot \otimes u_2$ for all $u_1, u_2 \in U$. Requiring (2.2) has very little effect on our results. Some of the proofs become easier, while others become somewhat more difficult, but the theorems still hold.

Note that we also do not require $\otimes$ to distribute over $\oplus$. The obvious way to state such a requirement is to require that $p \otimes (x \oplus y) = (p \otimes x) \oplus (p \otimes y)$. But this is not well-defined. The domain of $\otimes$ is $P \times U$ and the domain of $\oplus$ is $V \times V$. If $x, y \in U$, then $x \oplus y$, $p \otimes x$, and $p \otimes y$ are all well defined, but $x \oplus y$ may be an element of $V - U$, so $p \otimes (x \oplus y)$ may not be well defined. As we shall see, in cases where it the distributive property makes sense (for example, if $V = U$ or if $u_1 \oplus u_2 \in U$ for all $u_1, u_2 \in U$), then it actually does hold in many examples of interest.

Example 2.1 The *standard expectation domain*, which we denote $\mathbb{E}$, is $(\mathbb{R}, [0, 1], \mathbb{R}, +, \times)$, where the ordering on each domain is the standard order on the reals. This, of course, is the expectation

[3] Sometimes we use $\times$ to denote Cartesian product; the context will always make it clear whether this is the case.

domain which is used in defining most decision rules in the literature. It is clearly monotonic and has a + identity, namely 0. ∎

Example 2.2 Consider the expectation domain $E_2 = (\mathbb{R}, [0,1] \times [0,1], \mathbb{R} \times \mathbb{R}, \oplus, \otimes)$, where

- we use the standard order on the utility domain $\mathbb{R}$;
- the order $\preceq$ on the plausibility domain $[0,1] \times [0,1]$ is such that $(p_1, p_2) \preceq (q_1, q_2)$ iff $p_1 \leq q_1$ and $p_2 \leq q_2$;
- similarly, $(u_1, u_2) \precsim_V (v_1, v_2)$ iff $u_1 \leq v_1$ and $u_2 \leq v_2$;
- $\oplus$ is defined pointwise: $(u_1, u_2) \oplus (v_1, v_2) = (u_1 + v_1, u_2 + v_2)$;
- $\otimes$ is pointwise multiplication: $(p_1, p_2) \otimes u = (p_1 u, p_2 u)$.

We can view the utility domain $\mathbb{R}$ as a substructure of the valuation domain $\mathbb{R} \times \mathbb{R}$ by identifying the element $u \in \mathbb{R}$ with the pair (u, u). Note that the ordering on the plausibility domain and the ordering on the utility domain are both partial. E_2 is also monotonic, and has $(0,0)$ as the $\oplus$ identity. The distributive property (which makes sense here) is also easily seen to hold: $(p_1, p_2) \otimes (u_1 \oplus u_2) = ((p_1, p_2) \otimes u_1) \oplus ((p_1, p_2) \otimes u_2)$.

It turns out to also be of interest to consider the expectation domain E_2' which is defined just like E_2 except that the order $\precsim_V'$ on the valuation domain is defined by taking $(u_1, u_2) \precsim_V' (v_1, v_2)$ iff $\min(u_1, u_2) \leq \min(v_1, v_2)$. Note that this makes $\precsim_V'$ a total order. ∎

2.2 Decision Situations and Decision Problems

A *decision situation* (under uncertainty) describes the objective part of the circumstance that the DM faces (i.e., the part that is independent of the tastes and beliefs of the DM). We model a decision situation in a standard way, as a tuple $\mathcal{A} = (A, S, C)$, where

- S is the set of states of the world,
- C is the set of consequences, and
- A is a set of acts (i.e., a set of functions from S to C).

An act a is *simple* iff its range is finite. That is, a is simple if it has only finitely many consequences. Many works in the literature focus on simple acts (e.g., [Fishburn 1987]). We assume in this paper that A contains only simple acts; this means that we can define (generalized) expectation using finite sums, so we do not have to introduce infinite series or integration for arbitrary expectation domains. Note that all acts are guaranteed to be simple if either S or C is finite, although we do not assume that here.

A decision problem is essentially a decision situation together with information about the tastes and beliefs of the DM; that is, a decision problem is a decision situation together with the subjective part of the circumstance that faces the DM. Formally, a *(plausibilistic) decision problem* is a tuple $\mathcal{D} = (\mathcal{A}, E, \mathbf{u}, \mathrm{Pl})$, where

- $\mathcal{A} = (A, S, C)$ is a decision situation,
- $E = (U, P, V, \oplus, \otimes)$ is an expectation domain,
- $\mathbf{u} : C \to U$ is a utility function, and
- $\mathrm{Pl} : 2^S \to P$ is a plausibility measure.

We say that $\mathcal{D}$ is *monotonic* iff E is monotonic.

2.3 (Generalized) Expected Utility

Let $\mathcal{D} = ((A, S, C), E, \mathbf{u}, \mathrm{Pl})$ be a plausibilistic decision problem. Each $a \in A$ induces a *utility random variable* $\mathbf{u}_a : S \to U$ as follows: $\mathbf{u}_a(s) = \mathbf{u}(a(s))$. In the standard setting (where utilities are real-valued and Pl is a probability measure Pr), we can identify the expected utility of act a with the expected value of $\mathbf{u}_a$ with respect to Pr, computed in the standard way (where we use $\mathrm{ran}(f)$ to denote the range of a function f):

$$\mathbf{E}_{\mathrm{Pr}}(\mathbf{u}_a) = \sum_{x \in \mathrm{ran}(\mathbf{u}_a)} \mathrm{Pr}(\mathbf{u}_a^{-1}(x)) \times x.^4 \tag{2.3}$$

We can generalize (2.3) to an arbitrary expectation domain $E = (U, P, V, \oplus, \otimes)$ by replacing $+$, $\times$, and Pr by $\oplus$, $\otimes$, and Pl, respectively. This gives us

$$\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a) = \bigoplus_{x \in \mathrm{ran}(\mathbf{u}_a)} \mathrm{Pl}(\mathbf{u}_a^{-1}(x)) \otimes x. \tag{2.4}$$

We call (2.4) the *generalized EU* (GEU) of act a. Clearly (2.3) is a special case of (2.4).

In the probabilistic case, if all singleton sets are measurable with respect to Pr (i.e., in the domain of Pr), then

$$\mathbf{E}_{\mathrm{Pr}}(\mathbf{u}_a) = \sum_{s \in S} \mathrm{Pr}(s) \times \mathbf{u}_a(s). \tag{2.5}$$

The plausibilistic analogue of (2.5) is not necessarily equivalent to (2.4). A decision problem $((A, S, C), E, \mathbf{u}, \mathrm{Pl})$ is *additive* iff, for all $c \in C$ and nonempty $X, Y \subseteq S$ such that $X \cap Y = \emptyset$,

$$\mathrm{Pl}(X \cup Y) \otimes \mathbf{u}(c) = (\mathrm{Pl}(X) \otimes \mathbf{u}(c)) \oplus (\mathrm{Pl}(Y) \otimes \mathbf{u}(c)).$$

Note that the notion of additivity we defined is a joint property of several components of a decision problem (i.e., $\oplus$, $\otimes$, $\mathbf{u}$, and Pl) instead of being a property of Pl alone. Additivity is exactly the requirement needed to make the analogue of (2.5) equivalent to (2.4). While decision problems involving probability are additive, those involving representations of uncertainty such as Dempster-Shafer belief functions or, more generally, Choquet capacities, are not, in general.

Example 2.3 For a decision problem $(\mathcal{A}, \mathbb{E}, \mathbf{u}, \mathrm{Pr})$, where $\mathbb{E}$ is the standard expectation domain and $\mathbf{u}$ is a real-valued utility function, GEU agrees with EU. ∎

Example 2.4 Consider the decision problem$(\mathcal{A}, E_2, \mathbf{u}, (\mathrm{Pr}_1, \mathrm{Pr}_2))$, where E_2 is the expectation domain described in Example 2.2 and $\mathbf{u}$ is a real-valued utility function. The pair $(\mathrm{Pr}_1, \mathrm{Pr}_2)$ of probability measures can be viewed as a single plausibility measure. If $\mathcal{A} = (A, S, C)$, then the plausibility of $X \subseteq S$ is a pair $(\mathrm{Pr}_1(X), \mathrm{Pr}_2(X))$. It is easy to check that

$$\mathbf{E}_{(\mathrm{Pr}_1,\mathrm{Pr}_2),E_2}(\mathbf{u}_a) = (\mathbf{E}_{\mathrm{Pr}_1}(\mathbf{u}_a), \mathbf{E}_{\mathrm{Pr}_2}(\mathbf{u}_a).$$

Moreover, $\mathbf{E}_{(\mathrm{Pr}_1,\mathrm{Pr}_2),E_2}(\mathbf{u}_a) \precsim_V \mathbf{E}_{(\mathrm{Pr}_1,\mathrm{Pr}_2),E_2}(\mathbf{u}_{a'})$ iff $\mathbf{E}_{\mathrm{Pr}_i}(\mathbf{u}_a) \leq \mathbf{E}_{\mathrm{Pr}_i}(\mathbf{u}_{a'})$ for $i = 1, 2$.

On the other hand, if we consider E_2', we still have $\mathbf{E}_{(\mathrm{Pr}_1,\mathrm{Pr}_2),E_2'}(\mathbf{u}_a) = (\mathbf{E}_{\mathrm{Pr}_1}(\mathbf{u}_a), \mathbf{E}_{\mathrm{Pr}_2}(\mathbf{u}_a)$, but now $\mathbf{E}_{(\mathrm{Pr}_1,\mathrm{Pr}_2),E_2'}(\mathbf{u}_a)\mathbf{E}_{(\mathrm{Pr}_1,\mathrm{Pr}_2),E_2'}(\mathbf{u}_{a'})$ iff $\min(\mathbf{E}_{\mathrm{Pr}_1}(\mathbf{u}_a, \mathbf{E}_{\mathrm{Pr}_2}(\mathbf{u}_a)) \leq \min(\mathbf{E}_{\mathrm{Pr}_1}(\mathbf{u}_{a'}, \mathbf{E}_{\mathrm{Pr}_2}(\mathbf{u}_{a'}))$.

We can think of the plausibility measure $(\mathrm{Pr}_1, \mathrm{Pr}_2)$ as describing a situation where the DM is unsure which of Pr_1 and Pr_2 is the "right" probability measure. In E_2, act a is considered at least

[4] If the domain of Pr is some nontrivial subalgebra of 2^S, then we must assume that $\mathbf{u}_a$ is a measurable function; that is, $\mathbf{u}_a^{-1}(x)$ is a measurable set for all $x \in \mathrm{ran}(\mathbf{u}_a)$.

as good as a' if it is at least as good no matter which of $\Pr_1$ and $\Pr_2$ describes the actual situation. In E'_2, a is at least as good as a' if the worst expected outcome of a (with respect to each of $\Pr_1$ and $\Pr_2$) is at least as good as the worst expected outcome of a'. Note how relatively different orderings of valuation domain can produce quite different ordering on acts, using GEU. ∎

Another example of the use of GEU can be found in the proof of Theorem 3.1. Although the construction does not correspond to any standard decision problem in the literature, it does show how flexible the approach is.

3 Representing Arbitrary Preference Relations

In this section, we show that every preference relation on acts has a GEU representation. GEU, like all decision rules, is formally a function from decision problems to preference relations on acts. Thus a GEU representation of a preference relation $\precsim_A$ on the acts in $\mathcal{A} = (A, \ldots)$ is a decision problem $\mathcal{D} = (\mathcal{A}, E, \mathbf{u}, \mathrm{Pl})$, where $E = (U, P, V, \oplus, \otimes)$, such that $a_1 \precsim_A a_2$ iff $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{a_1}) \precsim_V \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{a_2})$.

Theorem 3.1 *Every preference relation $\precsim_A$ has a GEU representation.*

Proof: Fix some $\mathcal{A} = (A, S, C)$ and $\precsim_A$. We want to construct a decision problem $\mathcal{D} = (\mathcal{A}, E, \mathbf{u}, \mathrm{Pl})$ such that $\mathrm{GEU}(\mathcal{D}) = \precsim_A$.

The idea is to let each consequence be its own utility and each set be its own plausibility, and define $\otimes$ and $\oplus$ such that each act is its own expected utility. For each $c \in C$, let a_c denote the constant act with the property that $a_c(s) = c$ for all $s \in S$. Let $E = (U, P, V, \oplus, \otimes)$ be defined as follows:

1. $U = (C, \precsim_C)$, where $c \precsim_C d$ iff $c = d$ or $a_c, a_d \in A$ and $a_c \precsim_A a_d$. (Note that Savage assumes that A contains all simple acts; in particular, A contains all constant acts. We do not assume that here.)

2. $P = (2^S, \subseteq)$.

3. $V = (2^{S \times C}, \precsim_V)$, where $x \precsim_V y$ iff $x = y$ or $x, y \in A$ and $x \precsim_A y$. (Note that set-theoretically a function is a set of ordered pairs, so $A \subseteq 2^{S \times C}$.)

4. $x \oplus y = x \cup y$ for $x, y \in V$.

5. $X \otimes c = X \times \{c\}$ for $X \in 2^S$ $(= P)$ and $c \in C$ $(= U)$.

We can identify $c \in C$ with $S \times \{c\}$ in V; with this identification, $(U, \precsim_U)$ is a substructure of $(V, \precsim_V)$ and $\top \otimes c = c$ for all $c \in U$ $(= C)$, as required. Furthermore, $\oplus$ is clearly associative and commutative, so E is indeed an expectation domain. Let $\mathcal{D} = (\mathcal{A}, E, \mathbf{u}, \mathrm{Pl})$, where $\mathbf{u}(c) = c$ and $\mathrm{Pl}(X) = X$. Note that

$$\begin{aligned} \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a) &= \bigoplus\nolimits_{x \in \mathrm{ran}(\mathbf{u}_a)} \mathrm{Pl}(\mathbf{u}_a^{-1}(x)) \otimes x \\ &= \bigoplus\nolimits_{c \in \mathrm{ran}(a)} \mathrm{Pl}(a^{-1}(c)) \otimes c \\ &= \{(s, c) \mid a(s) = c\} \\ &= a. \end{aligned}$$

That is, each act is its own expected utility; by the definition of $\precsim_V$, it is clear that $a \precsim_A b$ iff $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a) \precsim_V \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_b)$. Thus $\mathrm{GEU}(\mathcal{D}) = \precsim_A$, as desired. ∎

Note that, unlike most representation theorems, there is no uniqueness condition in Theorem 3.1. This is because, unlike most representation theorems, we do not assume that the expectation domain is $\mathbb{E}$, the standard expectation domain, and we do not assume that $\precsim_A$ satisfies any assumptions. So one reason for the lack of uniqueness in Theorem 3.1 is because we place no restriction whatsoever on $\precsim_A$. The other reason for the lack of uniqueness is that we consider arbitrary expectation domains instead of restricting ourselves to $\mathbb{E}$. Note that, even if $\precsim_A$ satisfies all of Savage's postulates, although there is a unique GEU representation of $\precsim_A$ using the standard expectation domain $\mathbb{E}$ and probability measures (this is essentially Savage's result), there is no unique GEU representation if we allow arbitrary expectation domains. In particular, the representation constructed in the proof of Theorem 3.1 is certainly distinct from the one Savage [1954] constructs.

While there is no unique GEU representation, the GEU representation we construct in the proof of Theorem 3.1 is canonical in the following sense. Fix a decision situation $\mathcal{A} = (A, S, C)$ and a preference relation $\precsim_A$. Suppose $\mathcal{D} = (\mathcal{A}, E, \mathbf{u}, \mathrm{Pl})$ is the decision problem constructed in the proof of Theorem 3.1 and let $\mathcal{D}_0 = (\mathcal{A}, E_0, \mathbf{u}_0, \mathrm{Pl}_0)$ be an arbitrary GEU representation of $\precsim_A$, where $E_0 = (U_0, P_0, V_0, \widehat{\oplus}, \widehat{\otimes})$. It is easy to check that

- for all $X, Y \subseteq S$, $\mathrm{Pl}(X) \preceq_P \mathrm{Pl}(Y)$ implies $\mathrm{Pl}_0(X) \preceq_{P_0} \mathrm{Pl}_0(Y)$,
- for all $c, d \in C$, $\mathbf{u}(c) \precsim_U \mathbf{u}(d)$ implies $\mathbf{u}_0(c) \precsim_{U_0} \mathbf{u}_0(d)$, and
- for all $X_1, \ldots, X_n, Y_1, \ldots, Y_m \subseteq S$, for all $c_1, \ldots, c_n, d_1, \ldots, d_m \in C$,

$$\mathrm{Pl}(X_1) \otimes \mathbf{u}(c_1) \oplus \cdots \oplus \mathrm{Pl}(X_n) \otimes \mathbf{u}(c_n) \precsim_V \mathrm{Pl}(Y_1) \otimes \mathbf{u}(d_1) \oplus \cdots \oplus \mathrm{Pl}(Y_m) \otimes \mathbf{u}(d_m)$$

  implies

$$\mathrm{Pl}_0(X_1) \widehat{\otimes} \mathbf{u}_0(c_1) \widehat{\oplus} \cdots \widehat{\oplus} \mathrm{Pl}_0(X_n) \otimes \mathbf{u}_0(c_n) \precsim_{V_0} \mathrm{Pl}_0(Y_1) \widehat{\otimes} \mathbf{u}_0(d_1) \widehat{\oplus} \cdots \widehat{\oplus} \mathrm{Pl}_0(Y_m) \otimes \mathbf{u}_0(d_m).$$

Thus, the representation we construct is minimal, in the sense that we relate only what has to be related to satisfy the definition of representation. Moreover, the representation constructed in Theorem 3.1 is in fact additive, since if $X \cap Y = \emptyset$, then $(X \cup Y) \times \{c\} = (X \times \{c\}) \cup (Y \times \{c\})$, and has a $\oplus$ identity, namely $\emptyset$.

The expectation domain E constructed in the proof of Theorem 3.1 is not (necessarily) monotonic, since we certainly could have two acts a and b such that $a \precsim_A b$, so $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a) \precsim_V \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_b)$, but there is some $x \in V$ such that $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a) \oplus x \not\precsim_V \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_b) \oplus x$. In fact, our construction has the property that two distinct expressions are unrelated unless they are both expected utility values. As the following corollary shows, it is not hard to modify the proof by extending $\precsim_V$ so as to make E monotonic.

Corollary 3.2 *Every preference relation has a monotonic additive GEU representation with a $\oplus$ identity.*

Proof: See the appendix. ∎

Corollary 3.2 shows that requirements like monotonicity, additivity, and having a $\oplus$ identity do not restrict the kind of preference relation that GEU can represent. But this means that these requirements do not by themselves prevent GEU from producing "strange" preference relations when it is applied as a decision rule. In the next section, we consider constraints on expectations domains that do force the preference relation produced by GEU to be arguably more reasonable.

Theorem 3.1 holds in large part because of the flexibility we have. Given a decision situation (A, S, C) and a preference relation $\precsim_A$ on A, we are able to construct an expectation domain and a relation $\precsim_V$ that is customized to capture the relation $\precsim_A$ on A. We do not need quite this much flexibility. We can strengthen Theorem 3.1 to show that for every decision situation $\mathcal{A} = (A, S, C)$, there exists an expectation domain $E_{\mathcal{A}}$ such that for all preference relations $\precsim_A$ on A, there exists a utility function $\mathbf{u}$ and plausibility measure Pl such that $\mathrm{GEU}((\mathcal{A}, E, \mathbf{u}, \mathrm{Pl}) = \precsim_A$. That is, given $\mathcal{A}$, we can fix the expectation domain once and for all, rather than taking a different expectation structure (more precisely, a different order $\precsim_V$ on the valuation domain) for each preference relation $\precsim_A$. Indeed, we can even fix the plausibility measure once and for all as well.

Theorem 3.3 *Given a decision situation $\mathcal{A} = (A, S, C)$, there exists a monotonic, additive expectation domain E and a plausibility measure* Pl *on S such that, for every preference relation $\precsim_A$ on A, there exists a utility function $\mathbf{u}_{\precsim_A}$ on C and that $\precsim_A = \mathrm{GEU}(\mathcal{D})$, where $\mathcal{D} = (\mathcal{A}, E, \mathbf{u}_{\precsim_A}, \mathrm{Pl})$.*

Proof: Again, the argument proceeds by modifying the construction in Theorem 3.1. We leave details to the appendix. ∎

Theorems 3.1 and 3.3 depend (in part) on two features of our setup. The first is that, following Savage [1954], we took acts to be functions from states to consequences. This is not an entirely trivial assumption. In practice, different acts might produce the same consequences, depending on how the consequences are modeled. For example, suppose that Alice has a red umbrella and a blue umbrella (both in good condition). If the set of consequences is {"getting wet", "staying dry"}, then carrying the red umbrella will produce the same consequences as carrying the blue umbrella. Suppose instead that we have a *consequence function* $\mathbf{c} : A \times S \to C$ that takes an act a and a state s and gives the consequence of a in s. Of course, in this setting, two distinct acts a_1 and a_2 could induce the same function from states to consequences; that is, we might have $\mathbf{c}(a_1, s) = \mathbf{c}(a_2, s)$ for all $s \in S$. It is easy to see that if a_1 and a_2 induce the same function from states to consequences, then no matter what expectation domain, utility function, and plausibility measure we use, a_1 and a_2 will have the same expected utility. Thus, if $\precsim_A$ does not treat a_1 and a_2 the same way, then $\precsim_A$ has no GEU representation. (An analogue of Theorem 3.1 holds in this case: as long as two acts that induce the same function are treated the same way by $\precsim_A$, then $\precsim_A$ has a GEU representation.)

A second reason that we do not need consistency constraints on $\precsim_A$ is that we have placed no constraints on $\precsim_V$, and relatively few constraints on $\oplus$, $\otimes$, $\mathbf{u}$, and Pl. If, for example, we required $\precsim_V$ to be transitive, then we would also have to require that $\precsim_A$ be transitive. The lack of constraints on $\oplus$, $\otimes$, $\mathbf{u}$, and Pl is important because it gives us enough freedom to ensure that distinct acts have different expected utility. In the next section, we investigate what happens when we add more constraints.

4 Representing Savage's Postulates

Theorem 3.1 shows that GEU can represent any preference relation. We are typically interested in representing preference relations that satisfy certain constraints, or postulates. The goal of this section is to examine the effect of such constraints on the components that make up GEU. For definiteness, we focus on Savage's postulates.

A set $\mathcal{P}_e$ of axioms about (i.e., constraints on) plausibilistic decision problems *represents* a set of postulates $\mathcal{P}_r$ about decision situation and preference relation pairs *with respect to a collection of decision problems* Π iff for all $\mathcal{D} \in \Pi$,

$$\mathcal{D} = (\mathcal{A}, E, \mathbf{u}, \mathrm{Pl}) \text{ satisfies } \mathcal{P}_e \text{ iff } (\mathcal{A}, \mathrm{GEU}(\mathcal{D})) \text{ satisfies } \mathcal{P}_r.$$

Theorem 3.1 can be viewed as saying that the empty set of axioms represents the empty set of postulates with respect to the collection of all plausibilistic decision problems. Note that if $\mathcal{P}_e$ represents $\mathcal{P}_r$ with respect to Π_0 and $\Pi_1 \subseteq \Pi_0$, then $\mathcal{P}_e$ represents $\mathcal{P}_r$ with respect to $\mathcal{D}_1$ as well.

Before we present Savage's postulates, we first introduce some notation that will make the exposition more succinct. Suppose that $f : X \to Y$, $g : X \to Y$, and $Z \subseteq X$. Let $\langle f, Z, g \rangle$ denote the function h such that $h(x) = f(x)$ for all $x \in Z$ and $h(x) = g(x)$ for all $x \in \overline{Z}$. For example, if $X = Y = \mathbb{R}$ and $Z = \{x \mid x < 0\}$, then $\langle -x, Z, x \rangle$ is the absolute value function. In the intended application, the functions in question will be acts (i.e., functions from the set of states S to consequences C). So $a = \langle a_1, X, a_2 \rangle$ is the act such that $a(s) = a_1(s)$ for all $s \in X$ and $a(s) = a_2(s)$ for all $s \in \overline{X}$. For brevity, we identify the consequence $c \in C$ with the constant act a_c such that $a_c(s) = c$ for all $s \in S$. So for $c_1, c_2 \in C$, $\langle c_1, X, c_2 \rangle$ is the act with the property that $a(s) = c_1$ for all $s \in X$ and $a(s) = c_2$ for all $s \in \overline{X}$. Recall that $X_1, \ldots, X_n$ is a *partition* of Y iff the X_i's are nonempty and pairwise disjoint, and $\bigcup_i X_i = Y$.

Fix some decision situation (A, S, C). Readers familiar with [Savage 1954] will recall that Savage assumes that A consists of all possible functions from S to C, since the DM can be questioned about any pair of functions. (Though when Savage proves the main theorem in Chapter 5 of [Savage 1954], he restricts attention to acts that induce simple lotteries, since he essentially reduces his problem to the one already solved by von Neumann and Morgenstern [1947], and von Neumann and Morgenstern focused on simple lotteries.) This is a rather strong assumption. It means that the DM is required to have preferences on a rather large set of acts, many of which are not in his power to perform (and, indeed, many of which might be impossible to realize). Savage needs this assumption for his theorem. We do not need it for our results, although making this assumption simplifies the statement of the relevant axioms. As we have throughout this paper, in this section, we continue to allow A to be any nonempty subset of the set of all simple acts. The reader might wonder why we do not simply allow A to be the set of all simple acts, since we do not require $\precsim_A$ to be total. The point is that having A consist of all simple acts conceptually requires that the DM explicitly decide, for each pair of acts, whether they are related, and if so, how; if A is a subset of the set of all acts, then the DM does not have to express a preference between acts not in A.

It turns out that the statement of a number of our results is simpler if A consists of all simple acts. To facilitate the comparison of our results with the standard results from the literature, where it is typically assumed that A consists of all simple acts, we use brackets (i.e., "[" and "]") to delimit parts of the postulates that pertain to the general case in which A is an arbitrary nonempty subset of the set of all simple acts. So there are two versions of the postulates, one for the general case, which we refer to as the *general version*, and one for the special case (i.e., the case in which A is the set of all simple acts), which we refer to as the *special version*. The general version includes the bracketed statements while the special version does not. Typically, the statements inside the brackets turn unconditional assertions of the special version into implications whose antecedent says that the acts in question are in fact members of A. We recommend that the reader ignore the material inside the brackets on a first pass. Savage's first six postulates are given in Figure 1. It is easy to check that all the bracketed statements are trivially true if A is the set of all simple acts.

As is standard in the literature, we use "$a_1 \prec_A a_2$" to abbreviate "$a_1 \precsim_A a_2$ and $a_2 \not\precsim_A a_2$", and we use "$a_1 \sim_A a_2$" to abbreviate "$a_1 \precsim_A a_2$ and $a_2 \precsim_A a_1$". (Note that in general $\prec_A$ and $\sim_A$ are not necessarily transitive, since $\precsim_A$ is not necessarily transitive.) Recall that $X_1, \ldots, X_n$ is a *partition* of Y iff $\bigcup_i X_i = Y$ and for all $1 \leq i, j \leq n$ such that $i \neq j$, $X_i \neq \emptyset$ and $X_i \cap X_j = \emptyset$.

We now give a brief overview of the intuition behind the postulates and how Savage uses them. P1 is the standard necessary condition for representation by EU (and many of its generalizations), since $\mathbb{R}$ is a linear order; it basically says that $\precsim_A$ is a total preorder. Savage defines for each subset

P1. For all $a_1, a_2, a_3 \in A$,

(a) $a_1 \precsim_A a_2$ or $a_2 \precsim_A a_1$, and

(b) if $a_1 \precsim_A a_2$ and $a_2 \precsim_A a_3$, then $a_1 \precsim_A a_3$.

P2. For all $X \subseteq S$, $a_1, a_2, b_1, b_2 \in A$, $\big[$if $\langle a_i, X, b_j\rangle \in A$ for $i, j \in \{1,2\}$, then$\big]$

$$\langle a_1, X, b_1\rangle \precsim_A \langle a_2, X, b_1\rangle \text{ iff } \langle a_1, X, b_2\rangle \precsim_A \langle a_2, X, b_2\rangle.$$

P3. For all $X \subseteq S$, if there exist $a_1, a_2 \in A$ such that $\big[$there exists $b_0 \in A$ such that $\langle a_i, X, b_0\rangle \in A$ for $i \in \{1,2\}$, and$\big]$

for all $b \in A$, $\big[$if $\langle a_i, X, b\rangle \in A$ for $i \in \{1,2\}$, then$\big]$ $\langle a_1, X, b\rangle \prec_A \langle a_2, X, b\rangle$,

then for all $c_1, c_2 \in C$, $\big[$if $c_1, c_2 \in A$, then$\big]$ $c_1 \precsim_A c_2$ iff $\big[$there exists $b_0 \in A$ such that $\langle c_i, X, b_0\rangle \in A$ for $i \in \{1,2\}$, and$\big]$

for all $b \in A$, $\big[$if $\langle c_i, X, b\rangle \in A$ for $i \in \{1,2\}$, then$\big]$ $\langle c_1, X, b\rangle \precsim_A \langle c_2, X, b\rangle$.

P4. For all $X_1, X_2 \subseteq S$, $c_1, d_1, c_2, d_2 \in C$, if $\big[c_1, d_1, c_2, d_2 \in A,\big]$ $d_1 \prec_A c_1$ and $d_2 \prec_A c_2$, then $\big[$if $\langle c_i, X_j, d_i\rangle \in A$ for $i, j \in \{1,2\}$, then$\big]$

$$\langle c_1, X_1, d_1\rangle \precsim_A \langle c_1, X_2, d_1\rangle \text{ iff } \langle c_2, X_1, d_2\rangle \precsim_A \langle c_2, X_2, d_2\rangle.$$

P5. There exist $c_1, c_2 \in C$ such that $\big[c_1, c_2 \in A$ and$\big]$ $c_1 \prec_A c_2$.

P6. For all $a, b \in A$, $c \in C$, if $a \prec_A b$, then there exists a partition $Z_1, \ldots, Z_n$ of S, such that for all Z_i,

$\big[$if $\langle c, Z_i, a\rangle \in A$ then$\big]$ $\langle c, Z_i, a\rangle \prec_A b$ and

$\big[$if $\langle c, Z_i, b\rangle \in A$ then$\big]$ $a \prec_A \langle c, Z_i, b\rangle$.

Figure 1: Savage's Postulates

$X \subseteq S$ a conditional preference relation on acts as follows: $a_1 \precsim_A^X a_2$ iff [there exists $a \in A$ such that $\langle a_i, X, a\rangle \in A$ for $i \in \{1,2\}$ and]

$$\text{for all } a \in A, \text{ [if } \langle a_i, X, a\rangle \in A \text{ for } i \in \{1,2\}, \text{ then] } \langle a_1, X, a\rangle \precsim_A \langle a_2, X, a\rangle$$

(As in the statements of the postulates, we use brackets to delimit parts that are needed for the general version.) Intuitively, $a_1 \precsim_A^X a_2$ if when X occurs the DM would find a_2 at least as good as a_1. Note that $\precsim_A = \precsim_A^S$, so P1 guarantees that $\precsim_A^S$ is a total preorder. However, $\precsim_A^X$ is not necessarily a total preorder for all X, even if P1 holds—for this, we need P2.

P2 says that the way two acts are related depends only on where they differ; the part on which they agree can be ignored. Note that it follows from P2 that either

- for all $a \in A$, [if $\langle a_i, X, a\rangle \in A$ for $i \in \{1,2\}$, then] $\langle a_2, X, a\rangle \precsim_A \langle a_1, X, a\rangle$ or
- for all $a \in A$, $\langle a_2, X, a\rangle \not\precsim_A \langle a_1, X, a\rangle$.

Thus, in the presence of P2, $a_1 \precsim_A^X a_2$ iff

- for all $a \in A$, [if $\langle a_i, X, a\rangle \in A$ for $i \in \{1,2\}$, then] $\langle a_2, X, a\rangle \prec_A \langle a_1, X, a\rangle$ or
- for all $a \in A$, [if $\langle a_i, X, a\rangle \in A$ for $i \in \{1,2\}$, then] $\langle a_2, X, a\rangle \sim_A \langle a_1, X, a\rangle$.

Using $\precsim_A^X$, Savage defines what it means for a set to be null: a set X is *null* iff for all $a_1, a_2 \in A$, $a_1 \precsim_A^X a_2$ [iff $a_2 \precsim_A^X a_1$]. It is easy to check that if $\precsim_A^X$ is a total preorder then the general version and the special version are equivalent. Note that X is not null iff there exist $a_1, a_2 \in A$ such that $a_1 \prec_A^X a_2$. In other words, if X is not null, then the DM has some nontrivial preference if X occurs. P3 basically says that if X is not null, then $c_1 \precsim_A c_2$ iff $c_1 \precsim_A^X c_2$. That is, whenever the DM has some nontrivial preference, the preferences over consequences remain the same as the unconditional ones. Savage defines a relation $\precsim_S$ on events as follows: $X \precsim_S Y$ iff

$$\text{for all } c, d \in C, \text{ if } [c, d \in A,]\ d \prec_A c, \text{ [and } \langle c, X, d\rangle, \langle c, Y, d\rangle \in A,] \text{ then } \langle c, X, d\rangle \precsim_A \langle c, Y, d\rangle.$$

The intuition is that, given two consequences c and d such that $d \prec_A c$, the DM prefers a binary act that is more likely to yield c than d, according to her beliefs. This is very much in the spirit of arguments of de Finetti [1931]. P4, in the presence of P1–P3 and the assumption that A is the set of all simple acts, basically ensures that $\precsim_S$ is a total preorder. P5 says that S is not null. That is, the DM has some nontrivial (unconditional) preference. P1–P5 by themselves do not allow the construction of a unique EU representation (even if we assume that A is the set of all simple acts). However, with the assumption that A is the set of all simple acts, P1–P5 ensure that $\precsim_S$ is a qualitative probability. In order to obtain a unique EU representation we need P6, which says roughly that for all pairs of acts $a, b \in A$ and consequences $c \in C$, if $a \prec_A b$ then we can partition S into events such that the DM does not care if c were to happen in any element of the partition. Savage also has a seventh postulate, but it is relevant only for general (nonsimple) acts. Since we consider only simple acts, we omit it here.

It may seem that we should consider stronger versions of some postulates in the general case. For example, we might consider a version of P2 that says that if a_1, a_2, b_1, and b_2 are simple acts (not necessarily in A) such that $\langle a_i, X, b_j\rangle \in A$ for $1 \leq i, j \leq 2$, then $\langle a_1, X, b_1\rangle \precsim_A \langle a_2, X, b_1\rangle$ iff $\langle a_1, X, b_2\rangle \precsim_A \langle a_2, X, b_2\rangle$. Fortunately, it is not hard to show that the stronger version is equivalent to the version that we have stated here, where a_1, a_2, b_1, and b_2 are required to be in A, since if $\langle a, X, b\rangle \in A$, then in fact there exist some $a', b' \in A$ such that $\langle a', X, b'\rangle = \langle a, X, b\rangle$: just take $a' = b' = \langle a, X, b\rangle$. Thus it suffices in all the postulates to quantify over A instead of over all simple acts.

Given a decision problem $\mathcal{D} = ((A, S, C), E, \mathbf{u}, \mathrm{Pl})$ and $\emptyset \neq Z \subseteq S$, define the GEU of act a *restricted to* Z as follows:

$$\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a \restriction Z) = \bigoplus_{x \in \mathbf{u}_a(Z)} \mathrm{Pl}(\mathbf{u}_a^{-1}(x) \cap Z) \otimes x,$$

Note that $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a \restriction S) = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a)$. Suppose that $\mathcal{D} = ((A, S, C), E, \mathbf{u}, \mathrm{Pl})$ is additive. It is then easy to check that, for all nonempty proper subsets X of S,

$$\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a) = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a \restriction S) = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a \restriction X) \oplus \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a \restriction \overline{X}),$$

and, more generally, given a partition $X_1, \ldots, X_n$ of $Y \subseteq S$, we have that

$$\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a \restriction Y) = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a \restriction X_1) \oplus \cdots \oplus \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a \restriction X_n).$$

Also, it is easy to check that for all nonempty proper subsets X of S,

$$\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{\langle a_1, X, a_2 \rangle}) = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{a_1} \restriction X) \oplus \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{a_2} \restriction \overline{X}).$$

Note that while these statements are true for additive decision problems, they are not true in general.

Let $\mathcal{E}_{\mathcal{D}}(X) = \{\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a \restriction X) \mid a \in A\}$. (We omit the subscript $\mathcal{D}$ if it is clear from context.) Intuitively, $\mathcal{E}_{\mathcal{D}}(X)$ consists of all the expected utility values of acts in A restricted to X. To simplify the statement of one of the axioms, let

$$\langle\!\langle u, X, v \rangle\!\rangle = \begin{cases} u & \text{if } X = S, \\ v & \text{if } X = \emptyset, \\ \mathrm{Pl}(X) \otimes u \oplus \mathrm{Pl}(\overline{X}) \otimes v & \text{otherwise,} \end{cases}$$

where $u, v \in U$ and $X \subseteq S$. Note that $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{\langle c, X, d \rangle}) = \langle\!\langle \mathbf{u}(c), X, \mathbf{u}(d) \rangle\!\rangle$. The cases $X = \emptyset$ and $X = S$ must be treated specially, since we do not assume that $\perp \otimes u$ is the identity for $\oplus$. As with Savage's postulates, we use brackets to delimit parts needed for the general version. See Figure 2 for a list of the axioms.

A1 says that the expected utility values are linearly preordered; more specifically, A1a says that they are totally preordered and A1b says that the relation is transitive. Note that A1 does not say that the whole valuation domain is linearly preordered: that would be a sufficient but not a necessary condition for GEU($\mathcal{D}$) to satisfy P1. Since we want necessary and sufficient conditions for our representation results, some axioms apply only to expected utility values rather than to arbitrary elements of the valuation domain.

[There is a technical assumption that we need for some parts of the general version of our result. In general, it might be the case that $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a) \in \mathcal{E}(S)$, but $a \notin A$; this could happen if, even though $a \notin A$, there is some $b \in A$ such that $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a) = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_b)$. (Note that $\mathbf{u}_a$ is well defined whether or not $a \in A$, so it makes sense to talk about $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a)$ even if $a \notin A$.) We say that $\mathcal{D}$ is whole iff this does not happen; more precisely, $\mathcal{D} = ((A, S, C), E, \mathbf{u}, \mathrm{Pl})$ is *whole* iff for all simple acts $a \in C^S$, $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a) \in \mathcal{E}(S)$ implies $a \in A$. A decision problem whose set of acts is the set of all simple acts is whole, but that is not a necessary condition for a decision problem to be whole. In general, a decision problem $\mathcal{D} = ((A, S, C), E, \mathbf{u}, \mathrm{Pl})$, where $E = (U, P, V, \oplus, \otimes)$ is whole iff, for all $x \in V$, either every act with expected utility x is in A, or no act with expected utility x is in A.]

To simplify the statement of the theorem, let Π_{all} be the collection of all plausibilistic decision problems and let Π_{add} be the collection of additive decision problems. Also, let Π_0 be the collection of decision problems whose set of acts is the set of all simple acts [along with all decision problems that are whole]. Let

A1. For all $x, y, z \in \mathcal{E}(S)$,

(a) $x \precsim_V y$ or $y \precsim_V x$, and

(b) if $x \precsim_V y$ and $y \precsim_V z$, then $x \precsim_V z$.

A2. For all nonempty proper subsets X of S, $x_1, x_2 \in \mathcal{E}(X)$, $y_1, y_2 \in \mathcal{E}(\overline{X})$, $\llbracket$if $x_i \oplus y_j \in \mathcal{E}(S)$ for $i, j \in \{1, 2\}$, then$\rrbracket$

$$x_1 \oplus y_1 \precsim_V x_2 \oplus y_1 \text{ iff } x_1 \oplus y_2 \precsim_V x_2 \oplus y_2.$$

A3. For all nonempty proper subsets X of S, if there exist $x_1, x_2 \in \mathcal{E}(X)$ such that $\llbracket$there exists $y_0 \in \mathcal{E}(\overline{X})$ such that $x_i \oplus y_0 \in \mathcal{E}(S)$ for $i \in \{1, 2\}$, and$\rrbracket$

$$\text{for all } y \in \mathcal{E}(\overline{X}), \llbracket \text{if } x_i \oplus y \in \mathcal{E}(S) \text{ for } i \in \{1, 2\}, \text{ then} \rrbracket\ x_1 \oplus y \prec_V x_2 \oplus y,$$

then for all $u_1, u_2 \in \text{ran}(\mathbf{u})$, $\llbracket$if $u_1, u_2 \in \mathcal{E}(S)$, then$\rrbracket$ $u_1 \precsim_V u_2$ iff $\llbracket$there exists $y_0 \in \mathcal{E}(\overline{X})$ such that $\text{Pl}(X) \otimes u_i \oplus y_0 \in \mathcal{E}(S)$ for $i \in \{1, 2\}$, and$\rrbracket$

$$\text{for all } y \in \mathcal{E}(\overline{X}), \llbracket \text{if } \text{Pl}(X) \otimes u_i \oplus y \in \mathcal{E}(S) \text{ for } i \in \{1, 2\}, \text{ then} \rrbracket$$
$$\text{Pl}(X) \otimes u_1 \oplus y \precsim_V \text{Pl}(X) \otimes u_2 \oplus y.$$

A4. For all $X_1, X_2 \subseteq S$, $u_1, v_1, u_2, v_2 \in \text{ran}(\mathbf{u})$, if $\llbracket u_1, v_1, u_2, v_2 \in \mathcal{E}(S), \rrbracket$ $v_1 \prec_V u_1$ and $v_2 \prec_V u_2$, then $\llbracket$if $\langle\!\langle u_i, X_j, v_i \rangle\!\rangle \in \mathcal{E}(S)$ for $i, j \in \{1, 2\}$, then$\rrbracket$

$$\langle\!\langle u_1, X_1, v_1 \rangle\!\rangle \precsim_V \langle\!\langle u_1, X_2, v_1 \rangle\!\rangle \text{ iff } \langle\!\langle u_2, X_1, v_2 \rangle\!\rangle \precsim_V \langle\!\langle u_2, X_2, v_2 \rangle\!\rangle.$$

A5. There exist $u_1, u_2 \in \text{ran}(\mathbf{u})$ such that $\llbracket u_1, u_2 \in \mathcal{E}(S)$ and$\rrbracket$ $u_1 \prec_V u_2$.

A6. For all $x, y \in \mathcal{E}(S)$, $u \in \text{ran}(\mathbf{u})$, if $x \prec_V y$, then for all $a, b \in A$, $c \in C$, such that $\mathbf{E}_{\text{Pl},E}(\mathbf{u}_a) = x$, $\mathbf{E}_{\text{Pl},E}(\mathbf{u}_b) = y$, and $\mathbf{u}(c) = u$, there exists a partition $Z_1, \ldots, Z_n$ of S, such that x can be expressed as $x_1 \oplus \cdots \oplus x_n$ and y can be expressed as $y_1 \oplus \cdots \oplus y_n$, where $x_i = \mathbf{E}_{\text{Pl},E}(\mathbf{u}_a \restriction Z_i)$ and $y_i = \mathbf{E}_{\text{Pl},E}(\mathbf{u}_b \restriction Z_i)$ for $1 \leq i \leq n$, and for all $1 \leq k \leq n$,

$$\llbracket \text{if } \text{Pl}(Z_k) \otimes u \oplus \textstyle\bigoplus_{i \neq k} x_i \in \mathcal{E}(S) \text{ then} \rrbracket\ \text{Pl}(Z_k) \otimes u \oplus \textstyle\bigoplus_{i \neq k} x_i \prec_V y \text{ and}$$
$$\llbracket \text{if } \text{Pl}(Z_k) \otimes u \oplus \textstyle\bigoplus_{i \neq k} y_i \in \mathcal{E}(S) \text{ then} \rrbracket\ x \prec_V \text{Pl}(Z_k) \otimes u \oplus \textstyle\bigoplus_{i \neq k} y_i.$$

Figure 2: Axioms about Decision Problems

- $\Pi_{1a} = \Pi_{1b} = \Pi_5 = \Pi_{all}$, $\Pi_4 = \Pi_0$, and
- $\Pi_2 = \Pi_3 = \Pi_6 = \Pi_{add} \cap \Pi_0$.

Theorem 4.1 *For all* $i_1, \ldots, i_k \in \{1a, 1b, \ldots, 6\}$, $\{\mathrm{A}i_1, \ldots, \mathrm{A}i_k\}$ *represents* $\{\mathrm{P}i_1, \ldots, \mathrm{P}i_k\}$ *with respect to* $\Pi_{i_1} \cap \cdots \cap \Pi_{i_k}$.

Proof: See the appendix. ∎

Theorem 4.1 is a strong representation result. For example, if we are interested in capturing all of Savage's postulates but the requirement that $\precsim_A$ is a total preorder, and instead are willing to allow it to be a partial preorder (a situation explored by Lehmann [1996]), we simply need to drop the axiom A1b. Although we have focused here on Savage's postulates, it is straightforward to represent many of the other standard postulates considered in the decision theory literature in much the same way.

5 Conclusion

We have introduced GEU, a notion of generalized EU, and shown that GEU can (a) represent all preference relations on acts and (b) be customized to capture any subset of Savage's postulates. As we pointed out in the introduction, these results may be of particular interest to designers of software agents, who may want to deal with more general representations of tastes and beliefs than real-valued utilities and probabilities. If beliefs are represented using a plausibility measures and tastes are represented by a utility function that is not necessarily real-valued, the problem for the software designer is reduced to finding appropriate ways of combining plausibility and utility using $\oplus$ and $\otimes$, and finding an appropriate order on the resulting expressions. The results of this paper suggest that rationality postulates can be captured by choosing $\oplus$ and $\otimes$ so that they satisfy certain constraints. The results of [Chu and Halpern 2003] show that we lose no generality by using GEU to represent the decision making process; essentially all decision rules rules can be (ordinally) represented by GEU. Thus, the framework of expectation domains together with GEU provides a useful level of abstraction in which to study the general problem of decision making and rules for decision making and a useful conceptual framework for designing decision rules for software agents.

Acknowledgments

We thank Duncan Luce and Peter Wakker for useful comments on a draft of the paper.

A Proofs

Corollary 3.2 *Every preference relation has a monotonic additive GEU representation with a* $\oplus$ *identity.*

Proof: Fix some $\mathcal{A} = (A, S, C)$ and $\precsim_A$. Let $\mathcal{D}$ be as defined in the proof of Theorem 3.1, except that $x \precsim_V y$ iff $x = y$ or there exist $a, b \in A$ such that $a \precsim_A b$ and

1. $x = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a)$ and $y = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_b)$, or
2. $x = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a) \oplus z$ and $y = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_b) \oplus z$ for some $z \in V$.

Recall that $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a) = a$, so without case 2, we are back in the situation described in the proof of Theorem 3.1. Note that, by construction, the only way that $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a) \oplus z$ can be an expected utility value is if $z \subseteq \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a)$, since $\oplus = \cup$ and proper supersets of expected utility values cannot be expected utility values. Thus, if in case 2 both x and y are expected utility values, then we must in fact be in case 1; case 2 has an effect only when x and y are not both expected utility values. Thus, case 2 does not affect how pairs of expected utility values are related, so we still have that $\mathrm{GEU}(\mathcal{D}) = \precsim_A$.

To see that $\oplus$ is monotonic with respect to this definition of $\precsim_V$, suppose that $x \precsim_V y$. We need to show that $x \oplus z \precsim_V y \oplus z$. If $x = y$, then $x \oplus z = y \oplus z$, so the conclusion holds. Suppose that $x \neq y$. Then there exist some $a, b \in A$ such that $a \precsim_A b$ and either case 1 or case 2 holds. It is easy to see that in either case, $x \oplus z \precsim_V y \oplus z$ by case 2. Thus $\mathcal{D}$ is a monotonic representation of $\precsim_A$ (and, as we have already observed, $\mathcal{D}$ is additive). ∎

Theorem 3.3 *Given a decision situation $\mathcal{A} = (A, S, C)$, there exists a monotonic, additive expectation domain E and a plausibility measure* Pl *on S such that, for every preference relation $\precsim_A$ on A, there exists a utility function $\mathbf{u}_{\precsim_A}$ on C and that $\precsim_A = \mathrm{GEU}(\mathcal{D})$, where $\mathcal{D} = (\mathcal{A}, E, \mathbf{u}_{\precsim_A}, \mathrm{Pl})$.*

Proof: Let $\mathcal{P}(\mathcal{A})$ consist of all preference relations on A. We now modify the construction in Theorem 3.1 as follows:

1. $U = (C \times 2^{\mathcal{P}(\mathcal{A})}, \precsim_U)$, where $(c, X) \precsim_U (d, Y)$ iff $X = \{\precsim_A\} = Y$ for some $\precsim_A \in \mathcal{P}(\mathcal{A})$ and either $c = d$ or $a_c, a_d \in A$ and $a_c \precsim_A a_d$.
2. $P = (2^S, \subseteq)$.
3. $V = (2^{S \times C} \times 2^{\mathcal{P}(\mathcal{A})}, \precsim_V)$, where $(x, X) \precsim_V (y, Y)$ iff $X = \{\precsim_A\} = Y$ for some $\precsim_A \in \mathcal{P}(\mathcal{A})$ and either $x = y$ or $x, y \in A$ and $x \precsim_A y$.
4. $(x, X) \oplus (y, Y) = (x \cup y, X \cup Y)$.
5. $X \otimes (c, Y) = (X \times \{c\}, Y)$ for $X \subseteq S$, $c \in C$, and $Y \subseteq \mathcal{P}(\mathcal{A})$.

The same arguments as in the proof of Theorem 3.1, this construction gives an additive expectation domain. We can modify $\precsim_V$ as in Corollary 3.2 to make it monotonic. With a little more effort, we can further modify it so that there is a $\oplus$ identity; we omit details here.

Let $\mathrm{Pl}(X) = X$. Given a preference relation $\precsim_A$, define the utility function $\mathbf{u}_{\precsim_A}$ by taking $\mathbf{u}_{\precsim_A}(c) = (c, \{\precsim_A\})$. Again, the same arguments as those in Theorem 3.1 can be used to show that $\precsim_A = \mathrm{GEU}(\mathcal{D})$, where $\mathcal{D} = (\mathcal{A}, E, \mathbf{u}_{\precsim_A}, \mathrm{Pl})$. ∎

Theorem 4.1 *For all $i_1, \ldots, i_k \in \{1\mathrm{a}, 1\mathrm{b}, \ldots, 6\}$, $\{\mathrm{A}i_1, \ldots, \mathrm{A}i_k\}$ represents $\{\mathrm{P}i_1, \ldots, \mathrm{P}i_k\}$ with respect to $\Pi_{i_1} \cap \cdots \cap \Pi_{i_k}$.*

Proof: We first establish the result for singleton sets. Let $\mathcal{D} = (\mathcal{A}, E, \mathbf{u}, \mathrm{Pl})$, where $\mathcal{A} = (A, S, C)$, be an arbitrary decision problem. [As in the statements of the postulates and axioms, we will use brackets to delimit the parts of the proof that pertain to the conditional versions.]

- A1a represents P1a and with respect to $\Pi_{1\mathrm{a}}$ and A1b represents P1b with respect to $\Pi_{1\mathrm{b}}$.

  We do the case of A1a represents P1a with respect to $\Pi_{1\mathrm{a}}$ and leave the other case, which is completely analogous to the one we do, to the reader.

  Suppose that $\mathcal{D}$ satisfies A1a. We need to show that $(\mathcal{A}, \mathrm{GEU}(\mathcal{D}))$ satisfies P1a. Let $a_1, a_2, a_3 \in A$. Let $x_i = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{a_i})$; clearly $x_1, x_2, x_3 \in \mathcal{E}(S)$. Since $\mathcal{D}$ satisfies A1a, $x_1 \precsim_V x_2$

or $x_2 \precsim_V x_1$. In other words, $a_1 \precsim_{\mathrm{GEU}(\mathcal{D})} a_2$ or $a_2 \precsim_{\mathrm{GEU}(\mathcal{D})} a_1$. Thus $(\mathcal{A}, \mathrm{GEU}(\mathcal{D}))$ satisfies P1a. Now suppose that $(\mathcal{A}, \mathrm{GEU}(\mathcal{D}))$ satisfies P1a. We need to show that $\mathcal{D}$ satisfies A1a. Let $x_1, x_2, x_3 \in \mathcal{E}(S)$. Then there exist $a_1, a_2, a_3 \in A$ such that $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{a_i}) = x_i$. Since $(\mathcal{A}, \mathrm{GEU}(\mathcal{D}))$ satisfies P1a, $a_1 \precsim_{\mathrm{GEU}(\mathcal{D})} a_2$ or $a_2 \precsim_{\mathrm{GEU}(\mathcal{D})} a_1$. Thus $x_1 \precsim_V x_2$ or $x_2 \precsim_V x_1$. So $\mathcal{D}$ satisfies A1a.

- A2 represents P2 with respect to Π_2.

  Throughout this part of the proof, we assume that $\mathcal{D} \in \Pi_2$; in particular, we assume that $\mathcal{D}$ is additive and we will use this fact without further comment.

  Suppose that $\mathcal{D}$ satisfies A2. We need to show that $(\mathcal{A}, \mathrm{GEU}(\mathcal{D}))$ satisfies P2. Suppose that $X \subseteq S$ and $a_1, a_2, b_1, b_2 \in A$. [Suppose further that $\langle a_i, X, b_j \rangle \in A$.] We need to show that

  $$\langle a_1, X, b_1 \rangle \precsim_{\mathrm{GEU}(\mathcal{D})} \langle a_2, X, b_1 \rangle \text{ iff } \langle a_1, X, b_2 \rangle \precsim_{\mathrm{GEU}(\mathcal{D})} \langle a_2, X, b_2 \rangle.$$

  If $X = \emptyset$ or $X = S$, then the above is trivially true. So assume that X is a nonempty proper subset of S. Let $x_i = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{a_i} \upharpoonright X)$ and $y_j = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{b_j} \upharpoonright \overline{X})$. Clearly $x_i \in \mathcal{E}(X)$, $y_j \in \mathcal{E}(\overline{X})$, and $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{\langle a_i, X, b_j \rangle}) = x_i \oplus y_j$. [Furthermore, $x_i \oplus y_j \in \mathcal{E}(S)$, since $\langle a_i, X, b_j \rangle \in A$.] Since $\mathcal{D}$ satisfies A2, we have that

  $$x_1 \oplus y_1 \precsim_V x_2 \oplus y_1 \text{ iff } x_1 \oplus y_2 \precsim_V x_2 \oplus y_2,$$

  so

  $$\langle a_1, X, b_1 \rangle \precsim_{\mathrm{GEU}(\mathcal{D})} \langle a_2, X, b_1 \rangle \text{ iff } \langle a_1, X, b_2 \rangle \precsim_{\mathrm{GEU}(\mathcal{D})} \langle a_2, X, b_2 \rangle.$$

  Thus $(\mathcal{A}, \mathrm{GEU}(\mathcal{D}))$ satisfies P2.

  Suppose that $(\mathcal{A}, \mathrm{GEU}(\mathcal{D}))$ satisfies P2. We need to show that $\mathcal{D}$ satisfies A2. [We assume for this direction that $\mathcal{D}$ is whole.] Suppose that X is a nonempty proper subset of S, $x_1, x_2 \in \mathcal{E}(X)$, and $y_1, y_2 \in \mathcal{E}(\overline{X})$. [Suppose further that $x_i \oplus y_j \in \mathcal{E}(S)$.] We need to show that

  $$x_1 \oplus y_1 \precsim_V x_2 \oplus y_1 \text{ iff } x_1 \oplus y_2 \precsim_V x_2 \oplus y_2.$$

  Note that there exist $a_1, a_2, b_1, b_2 \in A$ such that $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{a_i} \upharpoonright X) = x_i$ and $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{b_i} \upharpoonright \overline{X}) = y_i$. Observe that $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{\langle a_i, X, b_j \rangle}) = x_i \oplus y_j$. [Since $x_i \oplus y_j \in \mathcal{E}(S)$ and $\mathcal{D}$ is whole, it follows that $\langle a_i, X, b_j \rangle \in A$.] Since $(\mathcal{A}, \mathrm{GEU}(\mathcal{D}))$ satisfies P2,

  $$\langle a_1, X, b_1 \rangle \precsim_{\mathrm{GEU}(\mathcal{D})} \langle a_2, X, b_1 \rangle \text{ iff } \langle a_1, X, b_2 \rangle \precsim_{\mathrm{GEU}(\mathcal{D})} \langle a_2, X, b_2 \rangle,$$

  so

  $$x_1 \oplus y_1 \precsim_V x_2 \oplus y_1 \text{ iff } x_1 \oplus y_2 \precsim_V x_2 \oplus y_2.$$

  Thus $\mathcal{D}$ satisfies A2.

- A3 represents P3 with respect to Π_3.

  Throughout this part of the proof, we assume that $\mathcal{D} \in \Pi_3$; in particular, we assume that $\mathcal{D}$ is additive and we will use this fact without further comment.

  For this part, we will prove a slightly stronger claim that actually has a shorter proof. Note that A3 and P3 are both implications. We show that $\mathcal{D}$ satisfies the antecedent (consequent) of A3 iff $(\mathcal{A}, \mathrm{GEU}(\mathcal{D}))$ satisfies the antecedent (consequent) of P3. [We assume that $\mathcal{D}$ is whole in these arguments.] This implies that $\mathcal{D}$ satisfies A3 iff $(\mathcal{A}, \mathrm{GEU}(\mathcal{D}))$ satisfies P3.

Note that P3 quantifies over all subsets of S while A3 quantifies over only nonempty proper subsets of S. It is easy to check that $\emptyset$ and S satisfy P3. (More precisely, $(\mathcal{A}, \mathrm{GEU}(\mathcal{D}))$ satisfies the instance of P3 in which X is instantiated with $\emptyset$ and the instance of P3 in which X is instantiated with S.) So for the rest of this part, we restrict our attention to nonempty proper subsets of S.

We begin by showing that $\mathcal{D}$ satisfies the antecedent of A3 iff $(\mathcal{A}, \mathrm{GEU}(\mathcal{D}))$ satisfies the antecedent of P3. Fix some X that is a nonempty proper subset of S. We need to show that there exist $x_1, x_2 \in \mathcal{E}(X)$ such that

1. [there exists $y_0 \in \mathcal{E}(\overline{X})$ such that $x_i \oplus y_0 \in \mathcal{E}(S)$ and] for all $y \in \mathcal{E}(\overline{X})$, [if $x_i \oplus y \in \mathcal{E}(S)$, then] $x_1 \oplus y \prec_V x_2 \oplus y$

iff there exist $a_1, a_2 \in A$ such that

2. [there exists $b_0 \in A$ such that $\langle a_i, X, b_0 \rangle \in A$ and] for all $b \in A$, [if $\langle a_i, X, b \rangle \in A$, then] $\langle a_1, X, b \rangle \prec_A \langle a_2, X, b \rangle$.

To see that 1 implies 2, suppose that $x_1, x_2 \in \mathcal{E}(X)$ satisfy 1. Since $x_1, x_2 \in \mathcal{E}(X)$, there exist $a_1, a_2 \in A$ such that $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{a_i} \restriction X) = x_i$. We show that a_1 and a_2 satisfy 2. [To see that the first conjunct is true, note that by 1 there exists $y_0 \in \mathcal{E}(\overline{X})$ such that $x_i \oplus y_0 \in \mathcal{E}(S)$; fix some such y_0. Note that there exists $b_0 \in A$ such that $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{b_0} \restriction \overline{X}) = y_0$; observe that $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{\langle a_i, X, b_0 \rangle}) = x_i \oplus y_0 \in \mathcal{E}(S)$. Since $\mathcal{D}$ is whole, $\langle a_i, X, b_0 \rangle \in A$. For the second conjunct, we proceed as follows.] Let $b \in A$ [be such that $\langle a_i, X, b \rangle \in A$]. We need to show that $\langle a_1, X, b \rangle \prec_A \langle a_2, X, b \rangle$. Let $y = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_b \restriction \overline{X})$. Note that $y \in \mathcal{E}(\overline{X})$ and $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{\langle a_i, X, b \rangle}) = x_i \oplus y$. [Furthermore, since $\langle a_i, X, b \rangle \in A$, $x_i \oplus y \in \mathcal{E}(S)$.] By 1, $x_1 \oplus y \prec_V x_2 \oplus y$; thus $\langle a_1, X, b \rangle \prec_A \langle a_2, X, b \rangle$.

To see that 2 implies 1, suppose that $a_1, a_2 \in A$ satisfy 2. Let $x_i = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{a_i} \restriction X)$; note that $x_i \in \mathcal{E}(X)$. We show that x_1 and x_2 satisfy 1. [To see that the first conjunct is true, note that by 2 there exists $b_0 \in A$ such that $\langle a_i, X, b_0 \rangle \in A$; fix some such b_0. Let $y_0 = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{b_0} \restriction \overline{X})$. Note that $x_i \oplus y_0 = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{\langle a_i, X, b \rangle})$, so $y_0 \in \mathcal{E}(\overline{X})$ and $x_i \oplus y_0 \in \mathcal{E}(S)$. For the second conjunct, we proceed as follows.] Let $y \in \mathcal{E}(\overline{X})$ [be such that $x_i \oplus y \in \mathcal{E}(S)$]. We need to show that $x_1 \oplus y \prec_V x_2 \oplus y$. Note that there exists $b \in A$ such that $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_b \restriction \overline{X}) = y$; observe that $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{\langle a_i, X, b \rangle}) = x_i \oplus y$. [Furthermore, since $x_i \oplus y \in \mathcal{E}(S)$ and $\mathcal{D}$ is whole, $\langle a_i, X, b \rangle \in A$.] By 2, $\langle a_1, X, b \rangle \prec_A \langle a_2, X, b \rangle$; thus $x_1 \oplus y \prec_V x_2 \oplus y$.

We now show that $\mathcal{D}$ satisfies the consequent of A3 iff $(\mathcal{A}, \mathrm{GEU}(\mathcal{D}))$ satisfies the consequent of P3. We need to show that

3. for all $u_1, u_2 \in \mathrm{ran}(\mathbf{u})$, [if $u_1, u_2 \in \mathcal{E}(S)$, then] $u_1 \precsim_V u_2$ iff [there exists $y \in \mathcal{E}(\overline{X})$ such that $\mathrm{Pl}(X) \otimes u_i \oplus y \in \mathcal{E}(S)$ and] for all $y \in \mathcal{E}(\overline{X})$, [if $\mathrm{Pl}(X) \otimes u_i \oplus y \in \mathcal{E}(S)$, then] $\mathrm{Pl}(X) \otimes u_1 \oplus y \precsim_V \mathrm{Pl}(X) \otimes u_2 \oplus y$

iff

4. for all $c_1, c_2 \in C$, [if $c_1, c_2 \in A$, then] $c_1 \precsim_{\mathrm{GEU}(\mathcal{D})} c_2$ iff [there exists $b \in A$ such that $\langle c_i, X, b \rangle \in A$ and] for all $b \in A$, [if $\langle c_i, X, b \rangle \in A$, then] $\langle c_1, X, b \rangle \precsim_{\mathrm{GEU}(\mathcal{D})} \langle c_2, X, b \rangle$.

Suppose that 3 holds. We need to show that 4 holds. Fix some $c_1, c_2 \in C$ [such that $c_1, c_2 \in A$]. Let $u_i = \mathbf{u}(c_i)$. Then $u_i \in \mathrm{ran}(\mathbf{u})$ [and $u_1, u_2 \in \mathcal{E}(S)$]. Note that $c_1 \precsim_{\mathrm{GEU}(\mathcal{D})} c_2$ iff $u_1 \precsim_V u_2$. By 3, $u_1 \precsim_V u_2$ iff [there exists $y \in \mathcal{E}(\overline{X})$ such that $\mathrm{Pl}(X) \otimes u_i \oplus y \in \mathcal{E}(S)$ and] for all $y \in \mathcal{E}(\overline{X})$, [if $\mathrm{Pl}(X) \otimes u_i \oplus y \in \mathcal{E}(S)$, then] $\mathrm{Pl}(X) \otimes u_1 \oplus y \precsim_V \mathrm{Pl}(X) \otimes u_2 \oplus y$. [It is easy to check that there exists $y \in \mathcal{E}(\overline{X})$ such that $\mathrm{Pl}(X) \otimes u_i \oplus y \in \mathcal{E}(S)$ iff there exists $b \in A$ such that $\langle c_i, X, b\rangle \in A$; the "only if" part depends on the assumption that $\mathcal{D}$ is whole.] To see that 4 holds, fix some $b \in A$ [such that $\langle c_i, X, b\rangle \in A$]. We need to show that $\langle c_1, X, b\rangle \precsim_{\mathrm{GEU}(\mathcal{D})} \langle c_2, X, b\rangle$. Let $y = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_b \upharpoonright \overline{X})$; then $y \in \mathcal{E}(\overline{X})$ and $\mathrm{Pl}(X) \otimes u_i \oplus y = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{\langle c_i, X, b\rangle})$. [Since $\langle c_i, X, b\rangle \in A$, $\mathrm{Pl}(X) \otimes u_i \oplus y \in \mathcal{E}(S)$]. By 3, $\mathrm{Pl}(X) \otimes u_1 \oplus y \precsim_V \mathrm{Pl}(X) \otimes u_2 \oplus y$; thus $\langle c_1, X, b\rangle \precsim_{\mathrm{GEU}(\mathcal{D})} \langle c_2, X, b\rangle$.

Now suppose that 4 holds. We need to show that 3 holds. Fix some $u_1, u_2 \in \mathrm{ran}(\mathbf{u})$ [such that $u_1, u_2 \in \mathcal{E}(S)$]. Then there exist some $c_1, c_2 \in C$ such that $\mathbf{u}(c_i) = u_i$ [and $c_1, c_2 \in A$]. Note that $u_1 \precsim_V u_2$ iff $c_1 \precsim_{\mathrm{GEU}(\mathcal{D})} c_2$. By 4, $c_1 \precsim_{\mathrm{GEU}(\mathcal{D})} c_2$ iff [there exists $b \in A$ such that $\langle c_i, X, b\rangle \in A$ and] for all $b \in A$, [if $\langle c_i, X, b\rangle \in A$, then] $\langle c_1, X, b\rangle \precsim_{\mathrm{GEU}(\mathcal{D})} \langle c_2, X, b\rangle$. [As before, it is easy to check that there exists $b \in A$ such that $\langle c_i, X, b\rangle \in A$ iff there exists $y \in \mathcal{E}(\overline{X})$ such that $\mathrm{Pl}(X) \otimes u_i \oplus y \in \mathcal{E}(S)$; now the "if" part depends on the assumption that $\mathcal{D}$ is whole.] To see that 3 holds, fix some $y \in \mathcal{E}(\overline{X})$ [such that $\mathrm{Pl}(X) \otimes u_i \oplus y \in \mathcal{E}(S)$]. We need to show that $\mathrm{Pl}(X) \otimes u_1 \oplus y \precsim_V \mathrm{Pl}(X) \otimes u_2 \oplus y$. Note that there exists some $b \in A$ such that $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_b \upharpoonright \overline{X}) = y$; observe that $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{\langle c_i, X, b\rangle}) = \mathrm{Pl}(X) \otimes u_i \oplus y$ [and $\langle c_i, X, b\rangle \in A$, since $\mathrm{Pl}(X) \otimes u_i \oplus y \in \mathcal{E}(S)$ and $\mathcal{D}$ is whole]. By 4, $\langle c_1, X, b\rangle \precsim_{\mathrm{GEU}(\mathcal{D})} \langle c_2, X, b\rangle$; thus $\mathrm{Pl}(X) \otimes u_1 \oplus y \precsim_V \mathrm{Pl}(X) \otimes u_2 \oplus y$.

- A4 represents P4 with respect to Π_4.

  Suppose that $\mathcal{D}$ satisfies A4. We need to show that $(\mathcal{A}, \mathrm{GEU}(\mathcal{D}))$ satisfies P4. Suppose that $X_1, X_2 \subseteq S$, $c_1, d_1, c_2, d_2 \in C$, [$c_1, d_1, c_2, d_2 \in A$, $\langle c_i, X_j, d_i\rangle \in A$,] $d_1 \prec_A c_1$, and $d_2 \prec_A c_2$. We need to show that

$$\langle c_1, X_1, d_1\rangle \precsim_{\mathrm{GEU}(\mathcal{D})} \langle c_1, X_2, d_1\rangle \text{ iff } \langle c_2, X_1, d_2\rangle \precsim_{\mathrm{GEU}(\mathcal{D})} \langle c_2, X_2, d_2\rangle.$$

  Let $u_i = \mathbf{u}(c_i)$ and $v_i = \mathbf{u}(d_i)$. Note that $u_1, v_1, u_2, v_2 \in \mathrm{ran}(\mathbf{u})$ [and $u_1, u_2, v_1, v_2 \in \mathcal{E}(S)$]. Also, $v_i \prec_V u_i$, since $d_i \prec_{\mathrm{GEU}(\mathcal{D})} c_i$. Note that $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{\langle c_i, X_j, d_i\rangle}) = \langle\!\langle u_i, X_j, v_i\rangle\!\rangle$. [Since $\langle c_i, X_j, d_i\rangle \in A$, $\langle\!\langle u_i, X_j, v_i\rangle\!\rangle \in \mathcal{E}(S)$.] Since $\mathcal{D}$ satisfies A4,

$$\langle\!\langle u_1, X_1, v_1\rangle\!\rangle \precsim_V \langle\!\langle u_1, X_2, v_1\rangle\!\rangle \text{ iff } \langle\!\langle u_2, X_1, v_2\rangle\!\rangle \precsim_V \langle\!\langle u_2, X_2, v_2\rangle\!\rangle,$$

  which means that

$$\langle c_1, X_1, d_1\rangle \precsim_{\mathrm{GEU}(\mathcal{D})} \langle c_1, X_2, d_1\rangle \text{ iff } \langle c_2, X_1, d_2\rangle \precsim_{\mathrm{GEU}(\mathcal{D})} \langle c_2, X_2, d_2\rangle.$$

  Thus, $(\mathcal{A}, \mathrm{GEU}(\mathcal{D}))$ satisfies P4.

  Now suppose that $(\mathcal{A}, \mathrm{GEU}(\mathcal{D}))$ satisfies P4. We need to show that $\mathcal{D}$ satisfies A4. [For this direction, we assume that $\mathcal{D}$ is whole.] Suppose that $X_1, X_2 \subseteq S$, $u_1, v_1, u_2, v_2 \in \mathrm{ran}(\mathbf{u})$, [$u_1, v_1, u_2, v_2 \in \mathcal{E}(S)$, $\langle\!\langle u_i, X_j, v_i\rangle\!\rangle \in \mathcal{E}(S)$,] $v_1 \prec_V u_1$, and $v_2 \prec_V u_2$. We need to show that

$$\langle\!\langle u_1, X_1, v_1\rangle\!\rangle \precsim_V \langle\!\langle u_1, X_2, v_1\rangle\!\rangle \text{ iff } \langle\!\langle u_2, X_1, v_2\rangle\!\rangle \precsim_V \langle\!\langle u_2, X_2, v_2\rangle\!\rangle.$$

  Let $c_1, d_1, c_2, d_2 \in C$ be such that [$c_1, d_1, c_2, d_2 \in A$,] $\mathbf{u}(c_i) = u_i$ and $\mathbf{u}(d_i) = v_i$. Then we see that $d_i \prec_{\mathrm{GEU}(\mathcal{D})} c_i$, since $v_i \prec_V u_i$. Note that $\langle\!\langle u_i, X_j, v_i\rangle\!\rangle = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{\langle c_i, X_j, d_i\rangle})$. [Since $\langle\!\langle u_i, X_j, v_i\rangle\!\rangle \in \mathcal{E}(S)$ and $\mathcal{D}$ is whole, $\langle c_i, X_j, d_i\rangle \in A$.] Since $(\mathcal{A}, \mathrm{GEU}(\mathcal{D}))$ satisfies P4,

$$\langle c_1, X_1, d_1\rangle \precsim_{\mathrm{GEU}(\mathcal{D})} \langle c_1, X_2, d_1\rangle \text{ iff } \langle c_2, X_1, d_2\rangle \precsim_{\mathrm{GEU}(\mathcal{D})} \langle c_2, X_2, d_2\rangle,$$

  which implies that

$$\langle\!\langle u_1, X_1, v_1\rangle\!\rangle \precsim_V \langle\!\langle u_1, X_2, v_1\rangle\!\rangle \text{ iff } \langle\!\langle u_2, X_1, v_2\rangle\!\rangle \precsim_V \langle\!\langle u_2, X_2, v_2\rangle\!\rangle.$$

  So $\mathcal{D}$ satisfies A4.

- A5 represents P5 with respect to Π_5.

  $\mathcal{D}$ satisfies A5 iff there exist $u_1, u_2 \in \operatorname{ran}(\mathbf{u})$ such that $[u_1, u_2 \in \mathcal{E}(S)$ and$]$ $u_1 \prec_V u_2$ iff there exist some $c_1, c_2 \in C$ such that $[c_1, c_2 \in A$ and$]$ $\mathbf{u}(c_1) \prec_V \mathbf{u}(c_2)$ iff there exist some $c_1, c_2 \in C$ such that $[c_1, c_2 \in A$ and$]$ $c_1 \prec_{\mathrm{GEU}(\mathcal{D})} c_2$ iff $(\mathcal{A}, \mathrm{GEU}(\mathcal{D}))$ satisfies P5.

- A6 represents P6 with respect to Π_6.

  Throughout this part of the proof, we assume that $\mathcal{D} \in \Pi_6$; in particular, we assume that $\mathcal{D}$ is additive and we will use this fact without further comment.

  Suppose that $\mathcal{D}$ satisfies A6. We need to show that $(\mathcal{A}, \mathrm{GEU}(\mathcal{D}))$ satisfies P6. Let $a, b \in A$ and $c \in C$. Suppose that $a \prec_{\mathrm{GEU}(\mathcal{D})} b$. We need to show that there exists a partition $Z_1, \ldots, Z_n$ of S, such that for all Z_i,

  1. $[$if $\langle c, Z_i, a\rangle \in A$ then$]$ $\langle c, Z_i, a\rangle \prec_{\mathrm{GEU}(\mathcal{D})} b$ and
  2. $[$if $\langle c, Z_i, b\rangle \in A$ then$]$ $a \prec_{\mathrm{GEU}(\mathcal{D})} \langle c, Z_i, b\rangle$.

  Let $x = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a)$, $y = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_b)$, and $u = \mathbf{u}(c)$. Then $x, y \in \mathcal{E}(S)$, $u \in \operatorname{ran}(\mathbf{u})$, $x \prec_V y$, $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a) = x$, $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_b) = y$, and $\mathbf{u}(c) = u$, so (by A6) there exists a partition $Z_1, \ldots, Z_n$ of S, such that x can be expressed as $x_1 \oplus \cdots \oplus x_n$ and y can be expressed as $y_1 \oplus \cdots \oplus y_n$, where $x_i = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a \restriction Z_i)$ and $y_i = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_b \restriction Z_i)$ for $1 \le i \le n$, and for all $1 \le k \le n$,

  3. $[$if $\mathrm{Pl}(Z_k) \otimes u \oplus \bigoplus_{i \neq k} x_i \in \mathcal{E}(S)$ then$]$ $\mathrm{Pl}(Z_k) \otimes u \oplus \bigoplus_{i \neq k} x_i \prec_V y$, and
  4. $[$if $\mathrm{Pl}(Z_k) \otimes u \oplus \bigoplus_{i \neq k} y_i \in \mathcal{E}(S)$ then$]$ $x \prec_V \mathrm{Pl}(Z_k) \otimes u \oplus \bigoplus_{i \neq k} y_i$.

  To see that 1 holds, note that $\mathrm{Pl}(Z_k) \otimes u \oplus \bigoplus_{i \neq k} x_i = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{\langle c, Z_i, a\rangle})$. $[$Suppose that $\langle c, Z_i, a\rangle \in A$; then $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{\langle c, Z_i, a\rangle}) \in \mathcal{E}(S)$.$]$ By 3, $\mathrm{Pl}(Z_k) \otimes u \oplus \bigoplus_{i \neq k} x_i \prec_V y$. Thus $\langle c, Z_i, a\rangle \prec_{\mathrm{GEU}(\mathcal{D})} b$ as desired. The argument that 2 holds is completely analogous (we use 4 instead of 3 to establish that 2 holds), and we leave it to the reader.

  Now suppose that $(\mathcal{A}, \mathrm{GEU}(\mathcal{D}))$ satisfies P6. We need to show that $\mathcal{D}$ satisfies A6. $[$For this direction we assume that $\mathcal{D}$ is whole.$]$ Let $x, y \in \mathcal{E}(S)$ and $u \in \operatorname{ran}(\mathbf{u})$. Suppose that $x \prec_V y$. Let $a, b \in A$ and $c \in C$ be such that $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a) = x$, $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_b) = y$, and $\mathbf{u}(c) = u$. We need to show that there exists a partition $Z_1, \ldots, Z_n$ of S, such that x can be expressed as $x_1 \oplus \cdots \oplus x_n$ and y can be expressed as $y_1 \oplus \cdots \oplus y_n$, where $x_i = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_a \restriction Z_i)$ and $y_i = \mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_b \restriction Z_i)$ for $1 \le i \le n$, and for all $1 \le k \le n$, 3 and 4 hold.

  Since $x \prec_V y$, $a \prec_{\mathrm{GEU}(\mathcal{D})} b$, so (by P6) there exists a partition $Z_1, \ldots, Z_n$ of S such that for all Z_i, 1 and 2 hold. To see that 4 holds, note that $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{\langle c, Z_i, b\rangle}) = \mathrm{Pl}(Z_k) \otimes u \oplus \bigoplus_{i \neq k} y_i$. $[$Suppose that $\mathbf{E}_{\mathrm{Pl},E}(\mathbf{u}_{\langle c, Z_i, b\rangle}) \in \mathcal{E}(S)$; then $\langle c, Z_i, b\rangle \in A$, since $\mathcal{D}$ is whole.$]$ By 2, $a \prec_A \langle c, Z_i, b\rangle$. Thus $x \prec_V \mathrm{Pl}(Z_k) \otimes u \oplus \bigoplus_{i \neq k} y_i$. The argument that 3 holds is completely analogous (we use 1 instead of 2 to establish that 3 holds), and we leave that to the reader.

So far we have shown that Ai represents Pi with respect to Π_i, for $i \in \{1\mathrm{a}, 1\mathrm{b}, \ldots, 6\}$. Let $i_1, \ldots, i_k \in \{1\mathrm{a}, 1\mathrm{b}, \ldots, 6\}$. Suppose that $\mathcal{D} \in \Pi_{i_1} \cap \cdots \cap \Pi_{i_k}$ and that $\mathcal{D}$ satisfies $\{\mathrm{A}i_1, \ldots, \mathrm{A}i_k\}$. Since $\mathcal{D} \in \Pi_{i_j}$ and $\mathcal{D}$ satisfies Ai_j, it follows that $(\mathcal{A}, \mathrm{GEU}(\mathcal{D}))$ satisfies Pi_j. Thus $(\mathcal{A}, \mathrm{GEU}(\mathcal{D}))$ satisfies $\{\mathrm{P}i_1, \ldots, \mathrm{P}i_k\}$. Conversely, if $\mathcal{D} \in \Pi_{i_1} \cap \cdots \cap \Pi_{i_k}$ and $(\mathcal{A}, \mathrm{GEU}(\mathcal{D}))$ satisfies $\{\mathrm{P}i_1, \ldots, \mathrm{P}i_k\}$, then $\mathcal{D}$ satisfies $\{\mathrm{A}i_1, \ldots, \mathrm{A}i_k\}$. Thus $\{\mathrm{A}i_1, \ldots, \mathrm{A}i_k\}$ represents $\{\mathrm{P}i_1, \ldots, \mathrm{P}i_k\}$ with respect to $\Pi_{i_1} \cap \cdots \cap \Pi_{i_k}$. ∎

References

Benvenuti, P. and R. Mesiar (2000). Integrals with respect to a general fuzzy measure. In M. Grabisch, T. Murofushi, and M. Sugeno (Eds.), *Fuzzy Measures and Applications—Theory and Applications*, pp. 205–232. Heidelberg: Physica Verlag.

Chu, F. and J. Y. Halpern (2003). Great expectations. Part II: Generalized expected utility as a universal decision rule. In *Proc. Eighteenth International Joint Conference on Artificial Intelligence (IJCAI '03)*.

de Finetti, B. (1931). Sul significato soggestivo del probabilità. *Fundamenta Mathematica 17*, 298–329.

Fishburn, P. C. (1987). Reconsiderations in the foundations of decision under uncertainty. *The Economic Journal 97*(388), 825–841.

Friedman, N. and J. Y. Halpern (1995). Plausibility measures: a user's guide. In *Proc. Eleventh Conference on Uncertainty in Artificial Intelligence (UAI '95)*, pp. 175–184.

Giang, P. H. and P. P. Shenoy (2001). A comparison of axiomatic approaches to qualitative decision making using possibility theory. In *Proc. Seventeenth Conference on Uncertainty in Artificial Intelligence (UAI 2001)*, pp. 162–170.

Gilboa, I. and D. Schmeidler (1989). Maxmin expected utility with a non-unique prior. *Journal of Mathematical Economics 18*, 141–153.

Gul, F. (1991). A theory of disappointment aversion. *Econometrica 59*(3), 667–686.

Kahneman, D. and A. Tversky (1979). Prospect theory: an analysis of decision under risk. *Econometrica 47*(2), 263–292.

Lehmann, D. (1996, August). Generalized qualitative probability: Savage revisited. In E. Horvitz and F. Jensen (Eds.), *Twelfth Conference on Uncertainty in Artificial Intelligence*, Portland, Oregon, pp. 381–388. Morgan Kaufmann.

Luce, R. D. (1990). Rational versus plausible accounting equivalences in preference judgments. *Psychological Science 1*, 225–234. Reprinted with minor changes in Ward Edwards (Ed.), *Utility Theories: Measurements and Applications*, pp. 187–206. Boston: Kluwer, 1992.

Luce, R. D. (2000). *Utility of Gains and Losses: Measurement-Theoretical and Experimental Approaches*. London: Lawrence Erlbaum.

Neumann, J. v. and O. Morgenstern (1947). *Theory of Games and Economic Behavior* (2nd ed.). Princeton, N.J.: Princeton University Press.

Quiggin, J. (1993). *Generalized Expected Utility Theory: The Rank-Dependent Expected Utility Model*. Boston: Kluwer.

Resnik, M. D. (1987). *Choices: An Introduction to Decision Theory*. Minneapolis: University of Minnesota Press.

Savage, L. J. (1954). *Foundations of Statistics*. New York: Wiley.

Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica 57*, 571–587.

Tversky, A. and D. Kahneman (1992). Advances in prospect theory: cumulative representation of uncertainty. *Journal of Risk and Uncertainty 5*, 297–323.

von Neumann, J. and O. Morgenstern (1947). *Theory of Games and Economic Behavior*. Princeton: Princeton University Press.

Yaari, M. E. (1987). The dual theory of choice under risk. *Econometrica 55*(1), 95–115.